arXiv:1703.01860v1 [cs.LO] 6 Mar 2017

The parameterized space complexity of model-checking bounded variable first-order logic *

Yijia Chen†
School of Computer Science
Fudan University, China
yijia.chen@cs.sjtu.edu.cn

Michael Elberfeld
Lehrstuhl für Informatik 7
RWTH Aachen, Germany
elberfeld@informatik.rwth-aachen.de

Moritz Müller‡
Kurt Gödel Research Center
University of Vienna, Austria
moritz.mueller@univie.ac.at

Abstract

The parameterized model-checking problem for a class of first-order sentences (queries) asks to decide whether a given sentence from the class holds true in a given relational structure (database); the parameter is the length of the sentence. In 1995 Vardi observed a polynomial time algorithm deciding the model-checking problem for queries with a bounded number of variables. We study its parameterized space complexity. For each bound on the quantifier alternation rank the problem becomes complete for the corresponding level of what we call the tree hierarchy, a hierarchy of parameterized complexity classes defined via space bounded alternating machines between parameterized logarithmic space and fixed-parameter tractable time. We observe that a parameterized logarithmic space model-checker for existential bounded variable queries would allow to improve Savitch's classical simulation of nondeterministic logarithmic space in deterministic space $O(\log^2)$. Further, we define a highly space efficient model-checker for queries with a bounded number of variables and bounded quantifier alternation rank. We study its optimality under the assumption that Savitch's theorem is optimal.

1 Introduction

The model-checking problem MC(FO) for first-order logic FO asks whether a given first-order sentence φ holds true in a given relational structure $\mathbf{A}$. This problem is of paramount importance

*Some of the results in this paper have appeared in the extended abstract[11]

†Supported by the National Natural Science Foundation of China (NSFC) under project 61373029 and the Sino-German Center for Research Promotion under the project CDZ 996.

‡Supported by the Austrian Science Fund (FWF) under project number P28699.

throughout computer science, and especially in database theory [36]. The problem is PSPACE-complete in general and even its restriction to primitive positive sentences and two-element structures stays NP-hard (cf. [36]). Hence neither syntactical nor structural restrictions seem to allow to get a handle on this problem. An important exception is Vardi's observation from 1995 that the natural bottom-up evaluation algorithm takes only polynomial time on sentences with a bounded number of variables [38]. Indeed, this algorithm runs in time $O(|\varphi| \cdot |\mathbf{A}|^s)$ on instances $(\mathbf{A}, \varphi)$ where φ is in FO^s, i.e. contains at most s many variables.

It has repeatedly been argued in the literature (see e.g. [36, 17]), that measuring computational resources needed to solve MC(FO) by functions in the length of the input only is unsatisfactory. It neglects the fact that in typical situations in database theory we are asked to evaluate a relatively short φ (the query) against a relatively large $\mathbf{A}$ (the database). Parameterized complexity theory measures computational resources by functions taking as an additional argument a *parameter* associated to the problem instance. For the parameterized model-checking problem $p\text{-MC}(\mathrm{FO})$ one takes the length $|\varphi|$ of φ as parameter and asks for fixed-parameter tractable restrictions of $p\text{-MC}(\mathrm{FO})$, i.e. restrictions decidable in *fpt time* $f(|\varphi|) \cdot |\mathbf{A}|^{O(1)}$ where $f : \mathbb{N} \to \mathbb{N}$ is an arbitrary computable function. This relaxed tractability notion allows for effective but inefficient formula manipulations and thereby the transfer of logical methods to the study of the problem (see [23, 22, 24] for surveys). For example, algorithms exploiting Gaifman's locality theorem [13, Theorem 2.5.1] solve $p\text{-MC}(\mathrm{FO})$ on structures of bounded local treewidth [19], or graphs excluding some minor [17] in fpt time, and on bounded degree graphs even in parameterized logarithmic space [15]. Parameterized logarithmic space, denoted by para-L, relaxes logarithmic space in much the same way as FPT relaxes polynomial time [4, 15].

Concerning restrictions on the syntactical side one can naturally stratify the problem into subproblems $p\text{-MC}(\Sigma_1)$, $p\text{-MC}(\Sigma_2), \ldots, p\text{-MC}(\mathrm{FO})$ according to quantifier alternation rank. These problems stay intractable in the sense that they are well known to be complete for the levels of the A-hierarchy $\mathrm{A}[1] \subseteq \mathrm{A}[2] \subseteq \cdots \subseteq \mathrm{AW}[*]$. The completeness stays true even in the presence of function symbols [16, p.201]. The research question is to understand for which sets Φ of first-order sentences, the problem $p\text{-MC}(\Phi)$ is tractable. In this introductory exposition let us restrict attention to decidable sets Φ in a (relational) vocabulary of bounded arity.

Today, the situation is well understood for sets Φ of primitive positive sentences, i.e. *conjunctive queries* in database terminology. Indeed, if Φ is a decidable such set of bounded arity, then, assuming $\mathrm{A}[1] \neq \mathrm{FPT}$, the model-checking probem $p\text{-MC}(\Phi)$ is fixed-parameter tractable if [12] and only if [25, 21] there is some constant $s \in \mathbb{N}$ such that every query in Φ is logically equivalent to a conjunctive query whose primal graph has treewidth at most $s + 1$. Now, these are precisely those conjunctive queries that can be written with at most s many variables [28, Lemma 5.2, Remark 5.3] (see also [12, Theorem 12] and [5, Theorem 4] for similar statements). In other words, if $p\text{-MC}(\Phi)$ is fixed-parameter tractable at all, then it can be decided by first preprocessing the query to one with a bounded number of variables and then run the natural evaluation algorithm. As pointed out in [5], this is a recurring paradigm. In fact, most sets Φ (not necessarily primitive positive) that have been identified to date to have a tractable $p\text{-MC}(\Phi)$, are contained in FO^s up to logical equivalence.

This gives special interest to the computational complexity of $p\text{-MC}(\mathrm{FO}^s)$. Moreover, already for $s = 2$ this problem encompasses problems of independent interest. For example, a directed

graph with two vertices named by constants s and t contains a path from s to t of length at most k if and only if it satisfies the sentence $\varphi_k(s) \in \mathrm{FO}^2$ (allowing constants) where $\varphi_k(x)$ is defined as follows:

$$\begin{array}{lcl} \varphi_0(x) & := & x{=}t, \\ \varphi_{k+1}(x) & := & \exists y\Big(\big(y{=}x \vee Exy\big) \wedge \exists x\big(x{=}y \wedge \varphi_k(x)\big)\Big) \end{array} \tag{1}$$

There is some recent work concerning the time complexity of p-$\mathrm{MC}(\mathrm{FO}^s)$ restricted to formulas *in prenex form* and binary (relational) vocabularies. This model-checking problem can be decided in time $f(|\varphi|) \cdot |\mathbf{A}|^{s-1}$ for some computable $f : \mathbb{N} \to \mathbb{N}$ while time $g(|\varphi|) \cdot |\mathbf{A}|^{s-1-\varepsilon}$ for any $\varepsilon > 0$ and any $g : \mathbb{N} \to \mathbb{N}$ would falsify the strong exponential time hypothesis [41], and, in fact, imply sub-quadratic time for a version of the Orthogonal Vectors problem [20].

Given the central importance of p-$\mathrm{MC}(\mathrm{FO}^s)$ it seems surprising that, to the best of our knowledge, its space complexity has not been thoroughly studied. Vardi [38] showed that for $s \geqslant 3$ the *classical* problem $\mathrm{MC}(\mathrm{FO}^s)$ is P-complete under logspace reductions. But, even assuming P $\neq$ L, this leaves open the possibility that p-$\mathrm{MC}(\mathrm{FO}^s)$ could be solved in *parameterized logarithmic space*, that is, (deterministic) space $f(|\varphi|) + O(\log |\mathbf{A}|)$ for some computable $f : \mathbb{N} \to \mathbb{N}$. The starting point of this paper is the question whether this is the case.[1]

As we shall see, both answers to this question would have breakthrough consequences in classical complexity theory. It is one of the central open questions or, in Lipton's words [30], "one of the biggest embarrassments" of complexity theory whether Savtich's upper bound $\mathrm{NL} \subseteq \mathrm{SPACE}(\log^2)$ can be improved. This is open since 1969 and there is a tangible possibility that *Savitch's theorem is optimal*, i.e.[2] $\mathrm{NL} \not\subseteq \mathrm{SPACE}(o(\log^2))$. Hemaspaandra et al. [26, Corollary 2.8] showed that Savitch's theorem is not optimal if there were problems of sublogarithmic density $o(\log n)$ and Turing hard for NL. See [33] for more information about this problem. We observe the following implications:

$$\mathrm{P} = \mathrm{L} \quad \Longrightarrow \quad p\text{-}\mathrm{MC}(\mathrm{FO}^s) \in \text{para-L} \quad \Longrightarrow \quad \text{Savitch's theorem is not optimal.} \tag{2}$$

These implications are not hard to verify. For the second, note that running the assumed model-checker on the sentences (1) solves the parameterized problem

p-$\mathrm{STCON}_{\leqslant}$	
Instance:	A directed graph $\mathbf{G}$, two vertices s and t, and $k \in \mathbb{N}$.
Parameter:	k.
Problem:	Is there a directed path of length at most k from s to t in $\mathbf{G}$?

On n-node graphs this requires only space $f(k)+O(\log n)$ for some computable f. Then there is a logspace computable unbounded function h such that the algorithm runs in logarithmic space on instances with $k \leqslant h(n)$. It is long known that such an algorithm implies that Savitch's theorem is not optimal (see e.g. [40]). We shall give the details of this argument, and in fact prove something stronger. More precisely the contributions of this paper are as follows.

[1] Relatedly, it has been asked by Flum and Grohe [15, Remark 26] whether p-$\mathrm{MC}(\mathrm{FO}^s)$ is in para-NL.

[2] By the space hierarchy theorem $\mathrm{SPACE}(s(n)) \neq \mathrm{SPACE}(\log^2 n)$ holds for the Turing machine model and for all (not necessarily space constructible) $s(n) \leqslant o(\log^2 n)$.

Our contributions. What is the complexity of p-MC(FO^s)? It is somewhere between para-L and FPT. Unfortunately not much is known about this complexity landscape, and only recently there is a slowly emerging picture [14, 7, 6, 11]. An obvious approach is to stratify the problem again according to quantifier alternation rank into subproblems p-MC(Σ_1^s), p-MC$(\Sigma_2^s), \ldots$. It turns out that there exist parameterized logarithmic space reductions from p-MC(FO^s) to p-MC(FO^2) (Corollary 4.3), and for each fixed t from p-MC(Σ_t^s) to p-MC(Σ_t^2) (Theorem 4.1). On the other hand we do not know whether one can reduce p-MC(Σ_{t+1}^2) to p-MC(Σ_t^s) for any s.

We give natural machine characterizations of the classes, denoted TREE$[t]$ and TREE$[*]$, of the problems reducible to p-MC(Σ_t^2) and p-MC(FO^2) respectively (Definition 3.1, Theorem 4.1). It turns out that TREE$[1]$ equals the class TREE introduced in [7], a parameterized analogue of LOGCFL. A parameterized analogue of NL, called PATH in [7], has been introduced in [14]. Unlike their classical counterparts [27, 37, 2], PATH and TREE are not known to be closed under complementation (this is a question from [7]). Now, our classes TREE$[2]$, TREE$[3], \ldots$ turn out to be just the natural alternation hierarchy above TREE. We arrive at what we view as a conceptual contribution of this paper, the *tree hierarchy*:

$$\text{para-L} \subseteq \text{PATH} \subseteq \text{TREE} = \text{TREE}[1] \subseteq \text{TREE}[2] \subseteq \cdots \subseteq \text{TREE}[*] \subseteq \text{FPT}. \tag{3}$$

We prove that the completeness of p-MC(Σ_t^s) for TREE$[t]$ holds also true in the presence of function symbols (Theorem 4.6). This situation within FPT is thus fully analogous to the situation with unboundedly many variables and the A-hierarchy. The proofs, however, are quite different. A further difference is that the tree hierarchy satisfies a collapse theorem (Corollary 4.2) like the polynomial hierarchy while such a theorem is unknown for the A-hierarchy.

The connection to the classical question whether Savitch's theorem is optimal reads, more precisely, as follows:

Theorem 1.1. *If Savitch's Theorem is optimal, then* PATH $\not\subseteq$ para-L.

Our opinion (expressed in [32]) is that the completeness results for the tree hierarchy can serve as the basis of a doable fine-structure theory of problems between para-L and FPT. However, they do not tell us much about the deterministic space complexity of p-MC(FO^s). A straightforward algorithm solves p-MC(FO^s) in space (cf. Lemma 6.1)

$$O(|\varphi| \cdot (\log|\varphi| + \log|\mathbf{A}|)).$$

We show that there might be not too much room for improvement:

Theorem 1.2. *Let f be an arbitrary function from $\mathbb{N}$ to $\mathbb{N}$. If Savitch's theorem is optimal, then p-MC(Σ_1^2) is not decidable in space $o(f(|\varphi|) \cdot \log|\mathbf{A}|)$.*

However, building on ideas of Ruzzo [34], we still manage to give a significant improvement for bounded quantifier alternation rank:

Theorem 1.3. *For all $s, t \in \mathbb{N}$ the problem p-MC(Σ_t^s) is decidable in space*

$$O(\log|\varphi| \cdot (\log|\varphi| + \log|\mathbf{A}|)).$$

It is unlikely that the bound on the quantifier alternation rank can be omitted in this statement. Indeed, it follows from Vardi's completeness result mentioned above that p-MC(FO^3) can not be decided in the displayed space unless P $\subseteq$ SPACE$(\log^2 n)$.

2 Preliminaries

For a natural number $n \in \mathbb{N}$ we set $[n] = \{1, \dots, n\}$ understanding $[0] = \emptyset$.

2.1 Structures

A *vocabulary* τ is a finite set of relation and function symbols. Relation and function symbols have an associated *arity*, a natural number. A τ*-structure* $\mathbf{A}$ consists in a finite nonempty set A, its *universe*, and for each r-ary relation symbol $R \in \tau$ an *interpretation* $R^{\mathbf{A}} \subseteq A^r$ and for each r-ary function symbol $f \in \tau$ an *interpretation* $f^{\mathbf{A}} : A^r \to A$. A *constant* is a function symbol c of arity 0. We identify its interpretation $c^{\mathbf{A}}$ with its unique value, an element of A.

A *directed graph* is an $\{E\}$-structure $\mathbf{G} = (G, E^{\mathbf{G}})$ for the binary relation symbol E. We refer to elements of G are *vertices* and to elements $(a, b) \in E^{\mathbf{G}}$ as *(directed) edges (from a to b)*. Note this allows $\mathbf{G}$ to have self-loops, i.e. edges from a to a. A *graph* is a directed graph $\mathbf{G} = (G, E^{\mathbf{G}})$ with irreflexive and symmetric $E^{\mathbf{G}}$. A *(directed) path* in a (directed) graph $\mathbf{G}$ is a sequence $(a_1, \dots, a_{k+1})$ of pairwise distinct vertices such that $(a_i, a_{i+1}) \in E^{\mathbf{G}}$ for all $i \in [k-1]$; the path is said to have *length* k and to be *from* a_0 *to* a_k. Note that there is a path of length 0 from every vertex to itself.

The *size* of a τ-structure $\mathbf{A}$ is

$$|\mathbf{A}| := |\tau| + |A| + \textstyle\sum_R |R^{\mathbf{A}}| \cdot \mathit{ar}(R) + \sum_f |A|^{\mathit{ar}(f)},$$

where R, f range over the relation and function symbols of τ respectively, and $\mathit{ar}(R)$, $\mathit{ar}(f)$ denote the arities of R, f respectively. For example, the size of a (directed) graph with n vertices and m edges is $O(n + m)$. Note that a reasonable binary ("sparse" or "list") encoding of $\mathbf{A}$ has length $O(|\mathbf{A}| \cdot \log |A|)$. The difference between the size as defined and the length of the binary encoding of a structure plays no role in this paper.

2.2 Formulas

Let τ be a vocabulary. A τ*-term* is a variable, or of the form $f t_1 \cdots t_r$ where f is an r-ary function symbol and $t_1, \dots, t_r$ are again τ-terms. *Atomic* τ-formulas, i.e. τ*-atoms*, have the form $t{=}t'$ or $R(t_1, \dots, t_r)$ where R is an r-ary relation symbol in τ and $t, t', t_1, \dots, t_r$ are τ-terms. General τ*-formulas* are built from atomic ones by $\wedge, \vee, \neg$ and universal and existential quantification $\forall x, \exists x$. The vocabulary τ is *relational* if it contains only relation symbols. For a tuple of variables $\bar{x} = (x_1, \dots, x_k)$ we write $\varphi = \varphi(\bar{x})$ to indicate that the free variables of φ are among $\{x_1, \dots, x_k\}$. If $\mathbf{A}$ is a τ-structure and $\bar{a} = (a_1, \dots, a_k) \in A^k$, then $\mathbf{A} \models \varphi(\bar{a})$ means that the assignment that maps x_i to a_i for $i \in [k]$ satisfies φ in $\mathbf{A}$. A *sentence* is a formula without free variables. The *size* $|\varphi|$ of a formula φ is the length of a reasonable binary encoding of it.

For $s \in \mathbb{N}$ let FO^s denote the class of (first-order) formulas *over a relational vocabulary* containing at most s variables (free or bound). For $t \in \mathbb{N}$ we define the classes Σ_t and Π_t as follows. Both Σ_0 and Π_0 are the class of quantifier free formulas; Σ_{t+1} (resp. Π_{t+1}) is the closure of Π_t (resp. Σ_t) under positive Boolean combinations (i.e. applying $\vee, \wedge$) and existential (resp. universal) quantification. We set

– $\Sigma_t^s := \mathrm{FO}^s \cap \Sigma_t$,

– $\Pi_t^s := \mathrm{FO}^s \cap \Pi_t$.

Example 2.1. The formulas (1) in the introduction do not qualify as FO^2 because they use constants s, t. They are evaluated in structures $\mathbf{G} = (G, E^{\mathbf{G}}, s^{\mathbf{G}}, t^{\mathbf{G}})$ where $(G, E^{\mathbf{G}})$ is a directed graph and $s^{\mathbf{G}}, t^{\mathbf{G}} \in G$ are two vertices. Let S, T be unary relation symbols and consider the $\{E, S, T\}$-structure $\mathbf{G}' = (G, E^{\mathbf{G}'}, S^{\mathbf{G}'}, T^{\mathbf{G}'})$ with $E^{\mathbf{G}'} := E^{\mathbf{G}}, S^{\mathbf{G}'} := \{s^{\mathbf{G}}\}$ and $T^{\mathbf{G}'} := \{t^{\mathbf{G}}\}$. Define formulas $\varphi'_k(x)$ as $\varphi_k(x)$ but with $\varphi'_0(x) := T(x)$. Then $\exists x(S(x) \wedge \varphi'_k(x))$ is in Σ_1^2 and

$$\mathbf{G} \models \varphi_k(s) \iff \mathbf{G}' \models \exists x(S(x) \wedge \varphi'_k(x)).$$

2.3 Parameterized complexity

A *(classical) problem* is a subset $Q \subseteq \{0,1\}^*$, where $\{0,1\}^*$ is the set of finite binary strings; the length of a binary string x is denoted by $|x|$. As model of computation we use Turing machines $\mathbb{A}$ with a (read-only) input tape and several worktapes. We shall consider Turing machines with nondeterminism and co-nondeterminism. For definiteness, let us agree that a *nondeterministic* Turing machine has special state $c_\exists, c_0, c_1$ and can nondeterministically move from state $c_\exists$ to state c_b with $b \in \{0,1\}$, and we say $\mathbb{A}$ *existentially guesses the bit* b. An *alternating* Turing machine additionally has a state $c_\forall$ allowing to *universally guess* a bit b. For a function $c : \{0,1\}^* \to \mathbb{N}$, the machine is said to *use* c *many (co-)nondeterministic bits* if for every $x \in \{0,1\}^*$ every run of $\mathbb{A}$ on x contains at most $c(x)$ many configurations with state $c_\exists$ (resp. $c_\forall$).

A *parameterized problem* is a pair (Q, κ) of a classical problem Q and a *parameterization* κ, i.e. a logarithmic space computable function $\kappa : \{0,1\}^* \to \mathbb{N}$ mapping any instance $x \in \{0,1\}^*$ to its *parameter* $\kappa(x) \in \mathbb{N}$.

We exemplify how we present parameterized problems. The *model-checking problem for* a class of first-order sentences Φ is the parameterized problem

p-MC(Φ)	
Instance:	A first-order sentence φ and a structure $\mathbf{A}$.
Parameter:	$\lvert\varphi\rvert$.
Problem:	$\varphi \in \Phi$ and $\mathbf{A} \models \varphi$?

More formally, e.g. p-MC(Φ) is the classical problem MC(Φ) containing all (binary strings encoding) pairs $(\varphi, \mathbf{A})$ with $\varphi \in \Phi$ and $\mathbf{A} \models \varphi$, together with a parameterization that maps binary strings encoding pairs of formulas and structures to the length of the binary string encoding the formula, and all other strings to 0.

The class FPT contains those parameterized problems (Q, κ) that can be decided in *fpt time (with respect to* κ*)*, i.e. in time $f(\kappa(x)) \cdot |x|^{O(1)}$ for some computable $f : \mathbb{N} \to \mathbb{N}$. The class para-L (para-NL) contains those parameterized problems (Q, κ) such that Q is decided (accepted) by some (nondeterministic) Turing machine $\mathbb{A}$ that runs in *parameterized logarithmic space (with respect to* κ*)*, i.e. in space $f(\kappa(x)) + O(\log |x|)$ for some computable function $f : \mathbb{N} \to \mathbb{N}$.

A *pl-reduction* from (Q, κ) to (Q', κ') is a reduction $R : \{0,1\}^* \to \{0,1\}^*$ from Q to Q' such that $\kappa'(R(x)) \leqslant f(\kappa(x))$ and $|R(x)| \leqslant f(\kappa(x)) \cdot |x|^{O(1)}$ for some computable $f : \mathbb{N} \to \mathbb{N}$, and R is *implicitly pl-computable*, that is, the following problem is in para-L:

p-BITGRAPH(R)	
Instance:	(x, i, b) with $x \in \{0,1\}^*$, $i \geqslant 1$, and $b \in \{0,1\}$.
Parameter:	$\kappa(x)$.
Problem:	Does $R(x)$ have length $\lvert R(x)\rvert \geq i$ and i-th bit b?

If there is such a reduction, (Q, κ) is *pl-reducible* to (Q', κ'), written $(Q, \kappa) \leqslant_{pl} (Q', \kappa')$. If also $(Q', \kappa) \leqslant_{pl} (Q, \kappa)$,the problems are *pl-equivalent*, written $(Q, \kappa) \equiv_{pl} (Q', \kappa')$.

It is routine to verify that $\leqslant_{pl}$ is transitive and $\equiv_{pl}$ an equivalence relation.

2.4 The classes PATH and TREE

The classes PATH and TREE capture the complexity of some natural problems in FPT up to pl-reductions. In fact, completeness for these classes yield two out of possibly four complexity degrees (with respect to pl-reductions) for model-checking classes of conjunctive queries. We refer to [7] for precise statements, and to [6] for an even finer classification. For our purpose it is sufficient to recall the definitions of these classes and one completeness result from [14].

The class PATH has been introduced in [14] and can be viewed as a parameterized analogue of NL. It contains those parameterized problems (Q, κ) such that Q is accepted by a nondeterministic Turing machine $\mathbb{A}$ which runs in parameterized logarithmic space, and uses $f(\kappa(x)) \cdot \log |x|$ many nondeterministic bits for some computable $f : \mathbb{N} \to \mathbb{N}$.

The following is from [14, Theorem 4]:

Theorem 2.2. *The parameterized problem* p-STCON$_{\leqslant}$ *is* PATH*-complete under pl-reductions.*

The class TREE has been introduced in [7] and can be viewed as a parameterized analogue of LOGCFL.[3] It contains the parameterized problems (Q, κ) such that Q is accepted by a nondeterministic Turing machine $\mathbb{A}$ which runs in parameterized logarithmic space, and for some computable $f : \mathbb{N} \to \mathbb{N}$ uses $f(\kappa(x)) \cdot \log |x|$ many nondeterministic bits and $f(\kappa(x))$ many co-nondeterministic bits.

3 The tree hierarchy

Following [7] we consider machines $\mathbb{A}$ with *mixed nondeterminism*. Additionally to the binary nondeterminism embodied in the states $c_\exists, c_\forall, c_0, c_1$ from Section 2.3 they use *jumps* explained as follows. Recall our Turing machines have an input tape. During a computation on an input x of length $n := |x| > 0$ the cells numbered 1 to n of the input tape contain the n bits of x. The machine has an *existential* and a *universal jump state* $j_\exists$ resp. $j_\forall$. A successor configuration of a configuration in a jump state is obtained by changing the state to the initial state and placing the input head on an arbitrary cell holding an input bit; the machine is said to *existentially resp. universally jump to* the cell.

[3] The analogy is based on the characterization of LOGCFL as logspace uniform SAC_1 [39].

Observe that the number of the cell to which the machine jumps can be computed in logarithmic space by moving the input head stepwise to the left. Intuitively, a jump should be thought as a guess of a number in $[n]$.

Acceptance is defined as usual for alternating machines. Call a configuration *universal* if it has state $j_\forall$ or $c_\forall$, and otherwise *existential*. The machine $\mathbb{A}$ accepts $x \in \{0,1\}^*$ if its initial configuration on x is *accepting*. The set of accepting configurations is the smallest set that contains all accepting halting configurations, that contains an existential configuration if it contains at least one of its successor configurations, and that contains a universal configuration if it contains all of its successor configurations.

Each run of $\mathbb{A}$ on some input x contains a subsequence of jump configurations (i.e. with state $j_\exists$ or $j_\forall$). For a natural number $t \geqslant 1$ the run is *t-alternating* if this subsequence consists in t blocks, the first consisting in existential configurations, the second in universal configurations, and so on. The machine $\mathbb{A}$ is *t-alternating* if every run of $\mathbb{A}$ on any input is t-alternating.

For $f : \{0,1\}^* \to \mathbb{N}$, we say $\mathbb{A}$ *uses f jumps (bits)* if for every $x \in \{0,1\}^*$ every run of $\mathbb{A}$ on x contains at most $f(x)$ many jump configurations (resp. configurations with state $c_\exists$ or $c_\forall$).

As for a more general notation, note that every run of $\mathbb{A}$ on x contains a (possibly empty) sequence of *nondeterministic configurations*, i.e. with state in $\{j_\exists, j_\forall, c_\exists, c_\forall\}$. The *nondeterminism type* of the run is the corresponding word over the alphabet $\{j_\exists, j_\forall, c_\exists, c_\forall\}$. For example, being $2t$-alternating means having nondeterminism type in $(\{j_\exists, c_\exists, c_\forall\}^*\{j_\forall, c_\exists, c_\forall\}^*)^t$. Here and below, we use regular expressions to denote languages over $\{j_\exists, j_\forall, c_\exists, c_\forall\}$.

Definition 3.1. A parameterized problem (Q, κ) is in $\textsc{Tree}[*]$ if there are a computable $f : \mathbb{N} \to \mathbb{N}$ and a machine $\mathbb{A}$ with mixed nondeterminism that accepts Q, runs in parameterized logarithmic space (with respect to κ) and uses $f \circ \kappa$ jumps and $f \circ \kappa$ bits. If additionally $\mathbb{A}$ is t-alternating for some $t \geqslant 1$, then (Q, κ) is in $\textsc{Tree}[t]$.

The definition of TREE$[t]$ is due to Hubie Chen (personal communication). The following two propositions are straightforward (cf. [7, Lemmas 4.5, 5.4]):

Proposition 3.2. *A parameterized problem (Q, κ) is in* PATH *if and only if there are a computable $f : \mathbb{N} \to \mathbb{N}$ and a 1-alternating machine with mixed nondeterminism that accepts Q, runs in parameterized logarithmic space (with respect to κ) and uses $f \circ \kappa$ jumps and 0 bits.*

Proposition 3.3. TREE $=$ TREE$[1]$.

Hence, all inclusions displayed in (3) in the introduction are trivial except possibly the last one TREE$[*] \subseteq$ FPT. We prove it in Corollary 4.4. It is likely to be strict:

Proposition 3.4.

1. para-NL $\subseteq$ TREE$[*]$ *if and only if* NL $\subseteq$ L.
2. FPT $\subseteq$ TREE$[*]$ *if and only if* P $\subseteq$ L.

Proof. We prove (1), the proof of (2) is similar. For the backward direction, note NL $\subseteq$ L implies para-NL $\subseteq$ para-L by general results of Flum and Grohe [15, Theorem 4, Proposition 8]. The

same results also imply the forward direction noting that TREE[$*$] is contained in the class XL. We include a direct argument: assume para-NL $\subseteq$ TREE[$*$], let Q be a classical problem which is NL-complete under logarithmic space reductions and let κ_0 be the parameterization which is constantly 0. Then $(Q, \kappa_0) \in$ para-NL $\subseteq$ TREE[$*$], so there are a function f and machine $\mathbb{A}$ with mixed nondeterminism that accepts Q and on input x uses $f(0) + O(\log |x|)$ space and $f(0)$ jumps and bits. Thus, Q can be decided in logarithmic space by simulating $\mathbb{A}$ for all possible outcomes of these constantly many jumps and bits. □

Remark 3.5. Elberfeld et al. [14] observed that para-NL $\subseteq$ PATH is equivalent to NL $\subseteq$ L.

The following technical lemma will prove useful in the next section.

Lemma 3.6 (Normalization)**.** *Let $t \geqslant 1$ and (Q, κ) be a parameterized problem.*

1. *$(Q, \kappa) \in$ TREE$[t]$ if and only if there are a computable $f : \mathbb{N} \to \mathbb{N}$ and a t-alternating machine $\mathbb{A}$ with mixed nondeterminism that accepts Q, runs in parameterized logarithmic space and such that for all $x \in \{0,1\}^*$ every run of $\mathbb{A}$ on x has nondeterminism type:*

$$\big((j_\exists c_\forall)^{f(\kappa(x))} (j_\forall c_\exists)^{f(\kappa(x))} \big)^{\lfloor t/2 \rfloor} (j_\exists c_\forall)^{f(\kappa(x)) \cdot (t \bmod 2)}. \tag{4}$$

2. *$(Q, \kappa) \in$ TREE$[*]$ if and only if there are a computable $f : \mathbb{N} \to \mathbb{N}$ and machine $\mathbb{A}$ with mixed nondeterminism that accepts Q, runs in parameterized logarithmic space and such that for all $x \in \{0,1\}^*$ every run of $\mathbb{A}$ on x has nondeterminism type:*

$$(j_\exists j_\forall)^{f(\kappa(x))}. \tag{5}$$

Proof. We only show (1), the proof of (2) is similar. The backward direction of (1) is obvious. To prove the forward direction, assume $(Q, \kappa) \in$ TREE$[t]$ and choose a parameterized logarithmic space t-alternating machine $\mathbb{A}$ with mixed nondeterminism accepting Q and a computable function $f : \mathbb{N} \to \mathbb{N}$ such that $\mathbb{A}$ on input $x \in \{0,1\}^*$ uses $f(\kappa(x))$ many jumps and bits.

Every run of $\mathbb{A}$ on x consists in at most t *blocks*. The first block is the sequence of configurations from the starting configuration until the first configuration in state $j_\forall$, the second block is the sequence starting from this configuration until the first following configuration in state $j_\exists$, and so on. We can assume that the first nondeterministic configuration is a jump configuration and thus has state $j_\exists$.

Define the machine $\mathbb{A}'$ as follows. On an input x of length at least 2, $\mathbb{A}'$ simulates $\mathbb{A}$ on x and additionally stores the parity of (the number of) the current block of the run. In an odd (even) block $\mathbb{A}'$ replaces $\mathbb{A}$'s existential (universal) guesses of a bit b by existential (universal) jumps. Namely, it first computes the number m of the cell currently read on the input tape, then performs an existential (universal) jump, then computes the parity b of the cell it jumps to, then moves the input head back to cell m and then continues the run of $\mathbb{A}$ with guessed bit b (i.e. in state c_b).

Then $\mathbb{A}'$ accepts Q and every run of $\mathbb{A}'$ on x does not have any configuration with state $c_\exists$ ($c_\forall$) in an odd (even) block. More precisely, and assuming $t = 3$ for notational simplicity: every run of $\mathbb{A}'$ on x has a nondeterminism type which is a prefix of a word in

$$j_\exists \{c_\forall, j_\exists\}^{\leqslant f(\kappa(x))} j_\forall \{c_\exists, j_\forall\}^{\leqslant f(\kappa(x))} j_\exists \{c_\forall, j_\exists\}^{\leqslant f(\kappa(x))}. \tag{6}$$

Since κ is computable in logarithmic space, the number $k := f(\kappa(x))$ can be computed in parameterized logarithmic space. Define a machine $\mathbb{A}''$ which on x first computes $k := f(\kappa(x))$ and then simulates $\mathbb{A}'$ keeping record of the nondeterminism type of the sofar simulated run. Moreover, $\mathbb{A}''$ uses the record to do appropriate dummy nondeterministic steps to ensure its nondeterminism type to be

$$j_\exists (c_\forall j_\exists)^{2k} c_\forall j_\forall (c_\exists j_\forall)^{2k} c_\exists j_\exists (c_\forall j_\exists)^{2k} c_\forall. \qquad \square$$

4 Model-checking problems and the tree hierarchy

4.1 Completeness results

It is easy to see that p-MC$(\mathrm{FO}^s) \in$ para-L for $s = 1$. We prove completeness results for $s \geqslant 2$.

Theorem 4.1. *Let $t \geqslant 1$ and $s \geqslant 2$. Then p-MC(Σ_t^s) is complete for* TREE$[t]$ *under pl-reductions.*

Proof. (Containment) We first show that p-MC$(\Sigma_t^s) \in$ TREE$[t]$. This is done by a straightforward algorithm $\mathbb{A}$ as follows. At any moment it keeps in memory a formula ψ and an assignment α to its free variables. For simplicity we assume that all negation symbols of ψ appear only in front of atoms.

In case ψ is an atom or an negated atom, $\mathbb{A}$ accepts if α satisfies ψ in $\mathbf{A}$ and rejects otherwise. If ψ is a disjunction $(\psi_0 \vee \psi_1)$ $\big($conjunction $(\psi_0 \wedge \psi_1)\big)$, the algorithm $\mathbb{A}$ existentially (universally) guesses a bit b and recurses replacing ψ by ψ_b and α by its restriction to the free variables in ψ_b. If ψ is $\exists x \psi_0$ (resp. $\forall x \psi_0$), then $\mathbb{A}$ makes an existential (resp. universal) jump to guess $a \in A$ and recurses replacing ψ by ψ_0 and α by the assignment which extends α mapping x to a. Here we assume x occurs freely in ψ_0 – otherwise $\mathbb{A}$ simply recurses on ψ_0 with α unchanged.

When started on a Σ_t^s-sentence φ and the empty assignment, the formulas occurring during the recursion are subformulas of φ and thus contain as most s variables, so each of the assignments computed during the recursion can be stored in space roughly $s \cdot \log |A|$.

(Hardness) We now show that p-MC(Σ_t^2) is hard for TREE$[t]$ under pl-reductions. Let $(Q, \kappa) \in$ TREE$[t]$ be given and choose a computable f and a t-alternating machine $\mathbb{B}$ with $f \circ \kappa$ jumps and $f \circ \kappa$ bits such that $\mathbb{B}$ accepts Q and runs in space $f(\kappa(x)) + O(\log |x|)$.

Given $x \in \{0,1\}^*$ compute an upper bound $s = f(\kappa(x)) + O(\log |x|)$ on the space needed by $\mathbb{B}$ on x. Since κ is computable in logarithmic space, such a number s can be computed in parameterized logarithmic space. We can assume that $\mathbb{B}$ on x always halts after at most $m = 2^{f(\kappa(x))} \cdot |x|^{O(1)}$ steps. Note the binary representation of m can be computed in parameterized logarithmic space. For two space s configurations c, c' of $\mathbb{B}$ on x, we say that $\mathbb{B}$ *reaches c' from c* if there is a length $\leqslant m$ computation of $\mathbb{B}$ leading from c to c' that does neither pass through a nondeterministic configuration nor through a configuration of space $> s$. We assume that $\mathbb{B}$ reaches a nondeterministic configuration from the initial configuration, i.e. the computation of $\mathbb{B}$ on x is not deterministic.

We define a structure $\mathbf{A}$ whose universe A comprises all $\big($length $O(s)$ binary codes of$\big)$ nondeterministic space s configurations of $\mathbb{B}$ on x. It interprets a binary relation symbol E, unary function symbols s_0, s_1 and unary relation symbols $S, F, J_\exists, J_\forall, C_\exists, C_\forall$ as follows.

A pair $(c, c') \in A^2$ is in $E^{\mathbf{A}}$ if there exists a successor configuration c'' of c such that $\mathbb{B}$ reaches c' from c''. The relation symbol S is interpreted by $S^{\mathbf{A}} = \{c_{\text{first}}\}$ where c_{first} is the (unique) first configuration in A reached by $\mathbb{B}$ from the initial configuration of $\mathbb{B}$ on x. The relation symbols $J_\exists, J_\forall, C_\exists$ and $C_\forall$ are interpreted by the sets of configurations in A with states $j_\exists, j_\forall, c_\exists$ and $c_\forall$ respectively. Obviously these sets partition A.

The relation symbol F is interpreted by the set $F^{\mathbf{A}}$ of those $c \in A$ such that one of the following holds:

- $c \in C_\exists^{\mathbf{A}} \cup J_\exists^{\mathbf{A}}$ and $\mathbb{B}$ reaches a space s accepting halting configuration from at least one successor configuration of c;
- $c \in C_\forall^{\mathbf{A}} \cup J_\forall^{\mathbf{A}}$ and $\mathbb{B}$ reaches a space s accepting halting configuration from all successor configurations of c.

Finally, the function symbols s_0 and s_1 are interpreted by any functions $s_0^{\mathbf{A}}, s_1^{\mathbf{A}} : A \to A$ such that for every $c \in C_\exists^{\mathbf{A}} \cup C_\forall^{\mathbf{A}}$ with $\{d \mid (c,d) \in E^{\mathbf{A}}\} \neq \emptyset$ we have:

$$\big\{s_0^{\mathbf{A}}(c), s_1^{\mathbf{A}}(c)\big\} = \big\{d \in A \;\big|\; (c,d) \in E^{\mathbf{A}}\big\}.$$

It is easy to check that $\mathbf{A}$ is computable from x in parameterized logarithmic space. For example, to check whether a given pair $(c, c') \in A^2$ is in $E^{\mathbf{A}}$ we simulate $\mathbb{B}$ starting from c for at most m steps; if the simulation wants to visit a configuration of space $> s$ or a nondeterministic configuration $\neq c'$, then we stop the simulation and reject.

For a word w of length $|w| \geqslant 1$ over the alphabet $\{j_\exists, j_\forall, c_\exists, c_\forall\}$ we define a formula $\varphi_w(x)$ with (free or bound) variables x, y as follows. We proceed by induction on the length $|w|$ of w.

If $|w| = 1$, define $\varphi_w(x) := F(x)$. For $|w| \geqslant 1$ define:

$$\begin{aligned}
\varphi_{c_\forall w}(x) &:= C_\forall(x) \wedge \big(\varphi_w(s_0(x)) \wedge \varphi_w(s_1(x))\big),\\
\varphi_{c_\exists w}(x) &:= C_\exists(x) \wedge \big(\varphi_w(s_0(x)) \vee \varphi_w(s_1(x))\big),\\
\varphi_{j_\exists w}(x) &:= J_\exists(x) \wedge \exists y\big(E(x,y) \wedge \exists x(x = y \wedge \varphi_w(x))\big),\\
\varphi_{j_\forall w}(x) &:= J_\forall(x) \wedge \forall y\big(\neg E(x,y) \vee \forall x(\neg x = y \vee \varphi_w(x))\big).
\end{aligned}$$

Let $|w| \geqslant 1$ and assume that $c \in A$ is a configuration such that every run of $\mathbb{B}$ on x starting at c has nondeterminism type w; then

$$c \text{ is accepting} \iff \mathbf{A} \models \varphi_w(c). \tag{7}$$

This follows by a straightforward induction on $|w|$. Now we look for $\mathbf{A}'$ and φ'_w with this property but in a relational vocabulary.

By the Normalization Lemma 3.6 we can assume that all runs of $\mathbb{B}$ on x have nondeterminism type w of the form (4). For such a w we observe that $\varphi_w(x) \in \Sigma_t^2$. Furthermore, the formula $\varphi_w(x)$ does not contain nested terms, in fact, all its atomic subformulas containing some function symbol are of the form $E(s_b(x), y)$, $J_\exists(s_b(x))$, or $J_\forall(s_b(x))$. For $b \in \{0,1\}$ we introduce binary relation symbols E_b and unary relation symbols $J_{\forall b}$ and $J_{\exists b}$, and then replace

the atomic subformulas $E(s_b(x), y)$, $J_\exists(s_b(x))$, $J_\forall(s_b(x))$ in $\varphi_w(x)$ by $E_b(x, y)$, $J_{\exists b}(x)$, $J_{\forall b}(x)$ respectively. This defines the formula $\varphi'_w(x)$. Since $\varphi_w(x) \in \Sigma^2_t$, also $\varphi'_w(x) \in \Sigma^2_t$.

To define $\mathbf{A}'$ we expand $\mathbf{A}$ by interpreting the new symbols by:

$$\begin{aligned} E_b^{\mathbf{A}'} &:= \big\{(c, d) \bigm| (s_b^{\mathbf{A}}(c), d) \in E^{\mathbf{A}}\big\}, \\ J_{\exists b}^{\mathbf{A}'} &:= \big\{c \bigm| s_b^{\mathbf{A}}(c) \in J_\exists^{\mathbf{A}}\big\}, \\ J_{\forall b}^{\mathbf{A}'} &:= \big\{c \bigm| s_b^{\mathbf{A}}(c) \in J_\forall^{\mathbf{A}}\big\}. \end{aligned}$$

We have for all $c \in A$:

$$\mathbf{A} \models \varphi_w(c) \iff \mathbf{A}' \models \varphi'_w(c).$$

As the assumption of (7) is satisfied for c_{first}, and c_{first} is accepting if and only if $\mathbb{B}$ accepts x, that is, if and only if $x \in Q$, we get

$$x \in Q \iff \mathbf{A}' \models \varphi'_w(c_{\text{first}}).$$

Setting $\psi := \exists x(S(x) \wedge \varphi'_w(x))$ we get a reduction as desired by mapping x to $(\psi, \mathbf{A}')$. □

4.2 Corollaries

The easy elimination of function symbols in the proof of Theorem 4.1 rests on the fact that it is applied only to formulas where function symbols are not nested. We shall treat the general case in the next subsection. Before that we draw some consequences of the last theorem and its proof.

Corollary 4.2. *Let $t' > t \geqslant 1$. If* $\textsc{Tree}[t]$ *is closed under complementation, then*

$$\textsc{Tree}[t'] = \textsc{Tree}[t].$$

Proof. It is sufficient to show this for $t' = t + 1$. Assume $\textsc{Tree}[t]$ is closed under complementation. By the previous theorem it suffices to show $p\text{-MC}(\Sigma^2_{t+1}) \in \textsc{Tree}[t]$, and we know $p\text{-MC}(\Sigma^2_t) \in \textsc{Tree}[t]$. By assumption we find a computable f and a t-alternating machine $\mathbb{A}'$ with $f \circ \kappa$ jumps and $f \circ \kappa$ bits that runs in parameterized logarithmic space and accepts the complement of $p\text{-MC}(\Sigma^2_t)$.

We describe a machine $\mathbb{B}$ accepting $p\text{-MC}(\Sigma^2_{t+1})$. On input $(\varphi, \mathbf{A})$ with a sentence $\varphi \in \Sigma^2_t$ we first eliminate all void universal quantifiers, i.e. replace all subformulas $\forall x\psi$ by ψ whenever x does not appear free in ψ. The machine $\mathbb{B}$ starts simulating the recursive algorithm $\mathbb{A}$ for $p\text{-MC}(\Sigma^2_{t+1})$ described in the previous proof. The simulation is stopped when $\mathbb{A}$ recurses to a subformula of φ of the form $\forall x\psi$. In this case, $\forall x\psi$ has at most one free variable y and the current assignment α maps it, say, to $a \in A$. The machine $\mathbb{B}$ simulates $\mathbb{A}'$ on $(\chi, \mathbf{A}')$ where $\mathbf{A}'$ expands $\mathbf{A}$ interpreting a new unary relation symbol C by $\{a\}$ and where χ is a Σ^2_t-sentence logically equivalent to $\neg\forall xy(C(y) \to \psi)$.

For the simulation of $\mathbb{A}'$ the machine $\mathbb{B}$ stores (ψ, a) and relies on the implicit pl-computability of $(\chi, \mathbf{A}')$ from $(\psi, a, \mathbf{A})$. □

Corollary 4.3. *Let $s \geqslant 2$. Then p-MC(FO^s) is complete for* TREE$[*]$ *under pl-reductions.*

Proof. That p-MC(FO^s) $\in$ TREE$[*]$ can be seen as in the proof of Theorem 4.1. Hardness of p-MC(FO^2) also follows as in this proof, but the argument is actually simpler: let $(Q, \kappa) \in$ TREE$[*]$ and choose a machine accepting it according to the Normalization Lemma 3.6 (2). Observe that the formula φ_w does not contain the function symbols s_0, s_1 for w as in (5). Hence the reduction can simply map $x \in \{0,1\}^*$ to $(\exists x(S(x) \wedge \varphi_w(x)), \mathbf{A}'')$ where $w = (j_\exists j_\forall)^{f(\kappa(x))}$ and $\mathbf{A}''$ is obtained from $\mathbf{A}$ by forgetting the interpretations of s_0, s_1. It follows from (7) that this defines a pl-reduction from (Q, κ) to p-MC(FO^2). □

Corollary 4.4. TREE$[*] \subseteq$ FPT.

Proof. By Corollary 4.3 and Vardi's result that the classical problem MC(FO^2) is in P [38]. □

As a further corollary we get that the parameterized analogue of Vardi's completeness result mentioned is likely false:

Corollary 4.5. *Unless* P $=$ L*, there is no $s \in \mathbb{N}$ such that p-*MC(FO^s) *is* FPT*-complete under pl-reductions.*

Proof. If p-MC(FO^s) is FPT-complete, then (we can assume $s \geqslant 2$ and thus) TREE$[*]$ $=$ FPT by Corollary 4.3. This implies P $=$ L by Proposition 3.4. □

4.3 Function symbols

Let func-Σ_t be defined as Σ_t except that function symbols are allowed. It is not hard to show (see e.g. [16, Example 8.55]) that p-MC(func-Σ_t) is equivalent to p-MC(Σ_t) under fpt-reductions (even pl-reductions). An analogous statement for the model-checking problems characterizing the classes of the W-hierarchy is not known. In fact, allowing function symbols gives problems complete for the presumably larger classes of the W$^{\text{func}}$-hierarchy. We refer to [10] for more information.

Let func-Σ_t^s be defined as Σ_t^s except that function symbols are allowed.

Theorem 4.6. *Let $s \geqslant 2$ and $t \geqslant 1$. Then p-*MC(Σ_t^s) $\equiv_{pl}$ p-MC(func-Σ_t^s).

Proof. By Theorem 4.1 it suffices to show p-MC(func-Σ_t^s) $\leqslant_{pl}$ p-MC(Σ_t^{s+3}). Let $\mathbf{A}$ be a structure of a vocabulary τ with $|A| \geqslant 2$. We define a relational vocabulary τ' which depends on τ only, and a τ'-structure $\mathbf{A}'$. The universe A' is the union

- of A.
- for each relation symbol $R \in \tau$ of arity r, of the set

  $$\big\{(a_1, \ldots, a_i) \;\big|\; 2 \leqslant i \leqslant r \text{ and } (a_1, \ldots, a_i, a_{i+1}, \ldots, a_r) \in R^{\mathbf{A}} \text{ for some } a_{i+1}, \ldots, a_r \in A\big\},$$

  i.e., the set of all 'partial' tuples that can be extended to some tuple in $R^{\mathbf{A}}$. Note that the size of this set is bounded by $(r-1) \cdot |R^{\mathbf{A}}|$.

– for each function symbol $f \in \tau$ of arity r, of the set

$$\bigcup_{2 \leqslant i \leqslant r} A^i.$$

We identify A^1 with A and it is therefore that the union $\bigcup_{2 \leqslant i \leqslant r} A^i$ starts at $i = 2$. Recall we assumed $|A| \geqslant 2$, so $\bigcup_{2 \leqslant i \leqslant r} A^i$ has size at most $|A|^{r+1}$. Altogether, $|A'| \leqslant |\mathbf{A}|^2$.

Now we define the vocabulary τ' and the τ'-structure $\mathbf{A}'$ in parallel:

– τ' contains a unary relation symbol U interpreted by the original universe of $\mathbf{A}$, that is,

$$U^{\mathbf{A}'} := A.$$

– For every constant symbol $c \in \tau$, the vocabulary τ' contains a unary relation symbol U_c which is interpreted as

$$U_c^{\mathbf{A}'} := \{c^{\mathbf{A}}\}.$$

– For every r-ary relation symbol $R \in \tau$, the vocabulary τ' contains a unary relation symbol U_R and we set

$$U_R^{\mathbf{A}'} := \{b \in A' \mid b = (a_1, \ldots, a_r) \text{ for some } (a_1, \ldots, a_r) \in R^{\mathbf{A}}\}.$$

- For every r-ary function symbol $f \in \tau$, the vocabulary τ' contains a binary relation symbol F_f and we define

$$F_f^{\mathbf{A}'} := \{((a_1, \ldots, a_r), a) \in A' \mid (a_1, \ldots, a_r) \in A^r \text{ and } f^{\mathbf{A}}(a_1, \ldots, a_r) = a\}.$$

– τ' contains a ternary function symbol R_e; the 'tuple *e*xtending relation' $R_e^{\mathbf{A}'}$ is defined by

$$R_e^{\mathbf{A}'} := \{(b, a, b') \in A' \times A \times A' \mid \text{there are } i \geqslant 1 \text{ and } a_1, \ldots, a_i \in A \text{ such that } b = (a_1, \ldots, a_i) \in A' \text{ and } b' = (a_1, \ldots, a_i, a) \in A'\},$$

Now, for any sentence $\varphi \in \text{func-}\Sigma_t^s$, we construct two sentences $\mathit{trans}_\varphi^\exists$ and $\mathit{trans}_\varphi^\forall$ such that

$$\begin{aligned} \mathbf{A} \models \varphi \quad &\iff \quad \mathbf{A}' \models \mathit{trans}_\varphi^\exists \\ &\iff \quad \mathbf{A}' \models \mathit{trans}_\varphi^\forall. \end{aligned} \tag{8}$$

We start by defining, for every term $m(\bar{x})$ with $\bar{x} = x_1, \ldots, x_s$, two formulas

$$\mathit{value}_m^\exists(x, \bar{x}) \quad \text{and} \quad \mathit{value}_m^\forall(x, \bar{x})$$

which respectively are in Σ_1^{s+3} and Π_1^{s+3} up to logical equivalence. Furthermore, for every $a \in A'$ and every $a_1, \ldots, a_s \in A$ it holds that

$$\begin{aligned} a = m^{\mathbf{A}}(a_1, \ldots, a_s) &\iff \mathbf{A}' \models \mathit{value}_m^\exists(a, a_1, \ldots, a_s) \\ &\iff \mathbf{A}' \models \mathit{value}_m^\forall(a, a_1, \ldots, a_s). \end{aligned}$$

If m is a variable x_i, then $\mathit{value}_m^{\exists} := \mathit{value}_m^{\forall} := x = x_i$. If m is a constant c, then $\mathit{value}_m^{\exists} := \mathit{value}_m^{\forall} := U_c(x)$. If m is the composed term $f(m_1, \ldots, m_r)$, then

$$\mathit{value}_m^{\exists}(x, \bar{x}) := \exists y \Big(\exists x \big(x = y \wedge \mathit{tuple}_{m_1,\ldots,m_r}(x, \bar{x}) \big) \wedge F_f(y, x) \Big),$$
$$\mathit{value}_m^{\forall}(x, \bar{x}) := \forall y \Big(\neg \exists x \big(x = y \wedge \mathit{tuple}_{m_1,\ldots,m_r}(x, \bar{x}) \big) \vee F_f(y, x) \Big),$$

where $\mathit{tuple}_{m_1,\ldots,m_r}$ is defined inductively on r as follows.

$$\mathit{tuple}_{m_1}(x, \bar{x}) := \mathit{value}_{m_1}^{\exists}(x, \bar{x}),$$
$$\mathit{tuple}_{m_1,\ldots,m_{i+1}}(x, \bar{x}) := \exists y \exists z \Big(R_e(y, z, x) \wedge \exists x \big(x = y \wedge \mathit{tuple}_{m_1,\ldots,m_i}(x, \bar{x}) \big) \wedge \exists x \big(x = z \wedge \mathit{value}_{m_i}^{\exists}(x, \bar{x}) \big) \Big).$$

It is easy to verify that $\mathit{tuple}_{m_1,\ldots,m_r} \in \Sigma_1^{s+3}$ and for every $b \in A'$ and every $a_1, \ldots, a_s \in A$

$$b = \Big(m_1^{\mathbf{A}}(a_1, \ldots, a_s), \ldots, m_r^{\mathbf{A}}(a_1, \ldots, a_s) \Big) \iff \mathbf{A}' \models \mathit{tuple}_{m_1,\ldots,m_r}(b, a_1, \ldots, a_s).$$

For formulas φ we define $\mathit{trans}_{\varphi}^{\exists}$ and $\mathit{trans}_{\varphi}^{\forall}$ by induction as follows:

$$\mathit{trans}_{m_1=m_2}^{\exists} := \exists x \big(\mathit{value}_{m_1}^{\exists}(x, \bar{x}) \wedge \mathit{value}_{m_2}^{\exists}(x, \bar{x}) \big),$$
$$\mathit{trans}_{m_1=m_2}^{\forall} := \forall x \big(\neg \mathit{value}_{m_1}^{\exists}(x, \bar{x}) \vee \mathit{value}_{m_2}^{\forall}(x, \bar{x}) \big),$$
$$\mathit{trans}_{Rm_1\ldots m_r}^{\exists} := \exists x \big(U_R(x) \wedge \mathit{tuple}_{m_1,\ldots,m_r}(x, \bar{x}) \big),$$
$$\mathit{trans}_{Rm_1\ldots m_r}^{\forall} := \forall x \big(\neg \mathit{tuple}_{m_1,\ldots,m_r}^{\exists}(x, \bar{x}) \vee U_R(x) \big).$$

If $\varphi = \neg\psi$, then $\mathit{trans}_{\varphi}^{\exists} := \neg \mathit{trans}_{\psi}^{\forall}$ and $\mathit{trans}_{\varphi}^{\forall} := \neg \mathit{trans}_{\psi}^{\exists}$. If $\varphi = (\psi_1 \vee \psi_2)$, then $\mathit{trans}_{\varphi}^{\exists} := \big(\mathit{trans}_{\psi_1}^{\exists} \vee \mathit{trans}_{\psi_2}^{\exists} \big)$ and $\mathit{trans}_{\varphi}^{\forall} := \big(\mathit{trans}_{\psi_1}^{\forall} \vee \mathit{trans}_{\psi_2}^{\forall} \big)$. The case for $(\psi_1 \wedge \psi_2)$ is similar. If $\varphi = \exists x_i \psi$, we define

$$\mathit{trans}_{\varphi}^{\exists} := \exists x_i \big(U_A(x_i) \wedge \mathit{trans}_{\psi}^{\exists} \big) \quad \text{and} \quad \mathit{trans}^{\forall}(\varphi) := \exists x_i \big(U_A(x_i) \wedge \mathit{trans}_{\psi}^{\forall} \big).$$

Similarly, for $\varphi = \forall x \psi$,

$$\mathit{trans}_{\varphi}^{\exists} := \forall x_i \big(\neg U_A(x_i) \vee \mathit{trans}_{\psi}^{\exists} \big) \quad \text{and} \quad \mathit{trans}_{\varphi}^{\forall} := \forall x_i \big(\neg U_A(x_i) \vee \mathit{trans}_{\psi}^{\forall} \big).$$

Now it is straightforward to show the equivalence (8). Moreover, if t is odd, then $\mathit{trans}_{\varphi}^{\exists}$ is equivalent to a Σ_t^{s+3}-sentence; and if t is even, then $\mathit{trans}_{\varphi}^{\forall}$ is equivalent to a Σ_t^{s+3}-sentence. For simplicity, we denote the corresponding Σ_t^{s+3}-sentence and Σ_t^{s+3}-sentence by $\mathit{trans}_{\varphi}^{\exists}$ and $\mathit{trans}_{\varphi}^{\forall}$ again. Therefore, for every structure $\mathbf{A}$ and $\varphi \in \text{func-}\Sigma_t^s$

$$\big(\mathbf{A}, \varphi \big) \mapsto \begin{cases} \big(\mathbf{A}', \mathit{trans}_{\varphi}^{\exists} \big) & \text{if } t \text{ is odd,} \\ \big(\mathbf{A}', \mathit{trans}_{\varphi}^{\forall} \big) & \text{if } t \text{ is even} \end{cases}$$

gives the desired pl-reduction from $p\text{-MC}(\text{func-}\Sigma_t^s)$ to $p\text{-MC}(\Sigma_t^{s+3})$. □

5 PATH and optimality of Savitch's theorem

Savitch's Theorem is a milestone result linking nondeterministic space to deterministic space.

STCON	
Instance:	A directed graph $\mathbf{G}$ and two vertices $s, t \in G$.
Problem:	Is there a (directed) path from s to t in $\mathbf{G}$?

Theorem 5.1 (Savitch [35])**.** $\text{STCON} \in \text{SPACE}(\log^2 n)$. *In particular,* $\text{NL} \subseteq \text{SPACE}(\log^2 n)$.

The second statement follows from the first via the trivial proposition:

Proposition 5.2. *Let* $s : \mathbb{N} \to \mathbb{N}$ *and assume* $\text{STCON} \in \text{SPACE}(s)$. *Then*

$$\text{NL} \subseteq \text{SPACE}\left(s\big(n^{O(1)}\big) + \log n\right).$$

In this section we prove a stronger version of Theorem 1.1 and Theorem 1.2. Additionally, we explain what a collapse of the parameterized classes PATH and para-L means in terms of classical complexity classes.

5.1 Proof of Theorem 1.1

Theorem 1.1 follows from the following stronger statement via Proposition 5.2.

Theorem 5.3. *Assume* $(\mathbf{G}, s, t, k) \overset{?}{\in} p\text{-STCON}_{\leqslant}$ *can be decided in deterministic space*

$$f(k) + O(\log |G|) \tag{9}$$

for some $f : \mathbb{N} \to \mathbb{N}$ *(not necessarily computable). Then* $\text{STCON} \in \text{SPACE}(o(\log^2 n))$.

Proof. Without loss of generality, assume that f in (9) satisfies $f(k) \geqslant k$ for every $k \in \mathbb{N}$. Let $\iota : \mathbb{N} \to \mathbb{N}$ be a non-decreasing and unbounded function such that for all $n \in \mathbb{N}$

$$f(\iota(n)) \leqslant \log n, \tag{10}$$

and hence

$$\iota(n) \leqslant \log n. \tag{11}$$

Note that we might not know how to compute $\iota(n)$.

Now let $\mathbf{G}$ be a directed graph, $s, t \in G$, $n := |G|$, and $k \geqslant 2$. We compute in space $O(\log k + \log n)$ the *minimum* $\ell := \ell(k) \in \mathbb{N}$ with

$$k^\ell \geqslant n - 1, \tag{12}$$

which implies

$$\ell \leqslant O\left(\frac{\log n}{\log k}\right). \tag{13}$$

Then we define a sequence of directed graphs $\big(\mathbf{G}_i^k\big)_{i\in\{0,\dots,\ell\}}$ with self-loops, as follows. For every $i \geqslant 0$ the vertex set G_i^k of $\mathbf{G}_i^k$ is the same as G, while for every $u, v \in G_i^k$ there is an edge from u to v in $\mathbf{G}_i^k$ if and only if there is a directed path from u to v in $\mathbf{G}$ of length at most k^i. In particular, $E^{\mathbf{G}_0^k}$ is the reflexive closure of $E^{\mathbf{G}}$. By (12)

$$\text{there is a path from } s \text{ to } t \text{ in } \mathbf{G} \iff \text{there is an edge from } s \text{ to } t \text{ in } \mathbf{G}_\ell^k. \tag{14}$$

Furthermore, for every $i \in [\ell]$ and $u, v \in G_i^k = G_{i-1}^k = G$

$$\begin{aligned}&\text{there is an edge from } u \text{ to } v \text{ in } \mathbf{G}_i^k\\ &\qquad\iff \text{there is a path from } u \text{ to } v \text{ in } \mathbf{G}_{i-1}^k \text{ of length at most } k.\end{aligned}$$

This can be decided by the following *recursive* algorithm:

$\mathbb{C}$
input: a directed graph $\mathbf{G}$, $k, i \in \mathbb{N}$, and $u, v \in G$
output: decide whether there is an edge in $\mathbf{G}_i^k$ from u to v.

1. **if** $i = 0$ **then** output whether $\big(u = v$ or $(u, v) \in E^{\mathbf{G}}\big)$ and return
2. simulate $\mathbb{A}$ on $\big(\mathbf{G}_{i-1}^k, u, v, k\big)$
3. **if** in the simulation of $\mathbb{A}$ queries "$(u', v') \in E^{\mathbf{G}_{i-1}^k}$?"
4. **then** call $\mathbb{C}(\mathbf{G}, k, i-1, u', v')$.

For every $k \geqslant 2$ let $\mathbb{C}^k$ be the algorithm which on every directed graph $\mathbf{G}$ and $s, t \in G$ first computes $\ell = \ell(k)$ as in (12) and then simulates $\mathbb{C}(\mathbf{G}, k, \ell, s, t)$. Thus, $\mathbb{C}^k$ decides whether there is a path from s to t in $\mathbf{G}$ by (14). We analyse its space complexity. First, the depth of the recursion tree is ℓ, as $\mathbb{C}^k$ recurses on $i = \ell, \ell-1, \dots, 0$. As usual, $\mathbb{C}^k$ has to maintain a stack of intermediate configurations for the simulations of

$$\mathbb{A}(\mathbf{G}_\ell^k, _, _, k), \mathbb{A}(\mathbf{G}_{\ell-1}^k, _, _, k), \dots, \mathbb{A}(\mathbf{G}_0^k, _, _, k).$$

For the simulation of each $\mathbb{A}(\mathbf{G}_i^k, _, _, k)$, the size of the configuration is linearly bounded by $f(k) + O(\log n)$ because of (9). Therefore, the total space required is

$$O\Big(\log k + \log n + \ell \cdot \big(f(k) + \log n\big)\Big) \leqslant O\left(\log k + \frac{f(k)\cdot \log n + \log^2 n}{\log k}\right) \tag{15}$$

by (13). As a consequence, in case $k = \iota(n)$, (10) and (11) imply that (15) is bounded by $o(\log^2 n)$. So if $\iota(n)$ would be computable, we could replace it with a logspace computable function, and then the result would follow. In particular, under the assumption PATH = para-L of Theorem 1.1, we can assume f is computable, and hence find a computable ι.

In order to circumvent the possible incomputability of $\iota(n)$ we adopt the strategy underlying Levin's optimal inverters [29, 8]. Namely, we simulate all the algorithms $\mathbb{C}^2, \mathbb{C}^3, \dots$ in a diagonal fashion, while slowly increasing the allowed space.

$\mathbb{S}$
input: a graph $\mathbf{G}$ and $s, t \in G$
output: decide whether there is a path in $\mathbf{G}$ from s to t.

1. $S \leftarrow 2$
2. **for all** $i = 2$ **to** S **do**
3. simulate $\mathbb{C}^i$ on $\big(\mathbf{G}, s, t\big)$ in *space* S
4. **if** the simulation accepts or rejects in space S
5. **then** accept or reject accordingly
6. $S \leftarrow S + 1$
7. goto 1.

Clearly $\mathbb{S}$ decides STCON. We prove that its space complexity is $o(\log^2 n)$ on every input graph $\mathbf{G}$ with $n := |G|$. To that end, let $k := \iota(n)$ and s^* be the space needed by $\mathbb{C}^k(\mathbf{G}, s, t)$. As argued before, we have $s^* \leqslant o(\log^2 n)$. Observe that $\mathbb{S}$ must halt no later than S reaches the value $\max(\iota(n), s^*)$, which is again bounded by $o(\log^2 n)$ due to (11). This concludes the proof. □

5.2 Proof of Theorem 1.2

Example 2.1 defines a pl-reduction from p-STCON$_\leqslant$ to p-MC(Σ_1^2) that maps an instance $(\mathbf{G}, s, t, k)$ of p-STCON$_\leqslant$ to an instance $(\varphi, \mathbf{A})$ of p-MC(Σ_1^2) such that $|\mathbf{A}| \leqslant O(|\mathbf{G}|)$ and $|\varphi| \leqslant c \cdot k$ for some constant $c \in \mathbb{N}$. Then, by the assumption and this reduction, there is an algorithm $\mathbb{A}$ which decides p-STCON$_\leqslant$ in space

$$g\big(f(c \cdot k) \cdot \log |\mathbf{G}|\big),$$

where $g : \mathbb{N} \to \mathbb{N}$ is a function with $\lim_{m\to\infty} g(m)/m = 0$. Here, we assume without loss of generality that the function f is non-decreasing, so $f(|\varphi|) \leqslant f(c \cdot k)$. Then for every $\ell \in \mathbb{N}$ there is an $m_\ell \in \mathbb{N}$ such that for every $m \geqslant m_\ell$ we have

$$\frac{g(m)}{m} \leqslant \frac{1}{\ell}.$$

Thus, for $\ell := f(c \cdot k)$ and $m := f(c \cdot k) \cdot \log |\mathbf{G}|$

$$g\big(f(c \cdot k) \cdot \log |\mathbf{G}|\big) = g\big(m\big) \leqslant \max_{m<m_\ell} g(m) + \frac{m}{\ell} = \max_{m<m_\ell} g(m) + \log |\mathbf{G}|.$$

Note $\max_{m<m_\ell} g(m)$ only depends on ℓ, and hence only on the parameter k. Consequently the space required by $\mathbb{A}$ can be bounded by

$$h(k) + \log |\mathbf{G}| \leqslant h(k) + O(\log |G|)$$

for an appropriate function $h : \mathbb{N} \to \mathbb{N}$. So the result follows from Theorem 5.3.

5.3 Bounded nondeterminism in logarithmic space

We close this section showing that the collapse of the parameterized classes para-L and PATH can be characterized as a collapse of L and NL restricted to 'arbitrarily few but non-trivially many' nondeterministic bits.

Definition 5.4. Let $c : \mathbb{N} \to \mathbb{N}$ be a function. The class NL$[c]$ contains all classical problems Q that are accepted by some nondeterministic Turing machine which uses $c(|x|)$ many nondeterministic bits and runs in logarithmic space.

Proposition 5.5. *The following are equivalent*

1. para-L = PATH.
2. *There exists a space-constructible*[4] *function* $c(n) \geqslant \omega(\log(n))$ *such that* NL$[c]$ = L.

Proof. To see that that (1) implies (2), assume para-L = PATH. Then there is a machine $\mathbb{A}$ deciding p-STCON$_\leqslant$ which on an instance $(\mathbf{G}, s, t, k)$ runs in space $f(k) + O(\log |G|)$ for some computable $f : \mathbb{N} \to \mathbb{N}$. We can assume that f is increasing and space-constructible. Then there is an unbounded, logarithmic space computable $\iota : \mathbb{N} \to \mathbb{N}$ such that $f(\iota(n)) \leqslant \log n$ for all $n \in \mathbb{N}$. Then

$$c(n) := \iota(n) \cdot \log n$$

is space-constructible and $c(n) \geqslant \omega(\log(n))$. We claim that NL$[c]$ = L.

Let $Q \in$ NL$[c]$ be given. Choose a machine $\mathbb{B}$ accepting Q that on input x uses at most $s = O(\log |x|)$ space and at most $\iota(|x|) \cdot \log |x|$ many nondeterministic bits. We assume there is at most one accepting space s configuration c_{acc} that can possibly appear in any run of $\mathbb{B}$ on x.

The digraph $\mathbf{G}$ has as vertices all space s configurations of $\mathbb{B}$ on x and an edge from u to v if there is a computation of $\mathbb{B}$ started at u leading to v which uses at most $\log |x|$ many nondeterministic bits and space at most s. Note this can be decided in logarithmic space by simulating $\mathbb{B}$ exhaustively for all possible outcomes of guesses.

Since ι is logarithmic space computable, the instance $\big(\mathbf{G}, c_{start}, c_{acc}, \iota(|x|)\big)$ of p-STCON$_\leqslant$ is implicitly logarithmic space computable from x; here, c_{start} is the starting configuration of $\mathbb{B}$ on x. This is a "yes" instance if and only if $x \in Q$. Given this input, $\mathbb{A}$ needs space

$$f(\iota(|x|)) + O(\log |G|) \leqslant O(\log |x|).$$

To see that (2) implies (1), assume $c(n) \geqslant \omega(\log(n))$ is space-constructible and NL$[c]$ = L. There is a logarithmic space computable, non-decreasing and unbounded function $\iota : \mathbb{N} \to \mathbb{N}$ such that $c(n) \geqslant \iota(n) \cdot \log n$ for all $n \in \mathbb{N}$. By Theorem 2.2 it suffices to show that p-STCON$_\leqslant$ can be decided in parameterized logarithmic space.

Define the classical problem

$$Q := \big\{(\mathbf{G}, s, t, k) \in \text{STCON}_\leqslant \mid k \leqslant \iota(|G|)\big\}.$$

[4]Recall, $c : \mathbb{N} \to \mathbb{N}$ is *space-constructible* if $c(n)$ can be computed from n in space $O(c(n))$.

Then $Q \in \mathrm{NL}[c]$ by a straightforward guess and check algorithm. By assumption, Q is decided by a logarithmic space algorithm $\mathbb{A}$. Then we solve p-STCON$_{\leqslant}$ as follows using some arbitrary "brute force" (deterministic) algorithm $\mathbb{B}$ deciding p-STCON$_{\leqslant}$.

Given an instance $(\mathbf{G}, s, t, k)$ we check whether $k \leqslant \iota(|G|)$; if this is the case we run $\mathbb{A}$ and otherwise $\mathbb{B}$. In the first case we consume only logarithmic space and in the second, the space is effectively bounded in k because the instance size is: if $k \geqslant \iota(|G|)$, then $|G| \leqslant f(k)$ where f is a computable function such that $f \circ \iota(n) \geqslant n$ for all $n \in \mathbb{N}$. □

Remark 5.6. There are similar characterizations of W[P] = FPT in [3, Theorem 3.8], EW[P] = EPT in [18, Theorem 4], and BPFPT = FPT in [31, Theorem 5.2].

We find it worthwhile to point out explicitly the following direct corollary.

Corollary 5.7. *If Savitch's Theorem is optimal, then* $\mathrm{L} \neq \mathrm{NL}[c]$ *for all space-constructible functions* $c(n) \geqslant \omega(\log n)$.

Proof. By Proposition 5.5 and Theorem 1.1. □

6 A deterministic model-checker

In this section we prove Theorem 1.3. This theorem claims for each $s \in \mathbb{N}$ the existence of an algorithm. Our algorithm is going to be uniform in s, so we formulate a more general statement formulating the space bound using the *width* $w(\varphi)$ of a formula φ. This is the maximal number of free variables in some subformula of φ.

Being interested in the dependence of the space complexity on the parameter $|\varphi|$, we formulate all results as statements about the classical model-checking problem $\mathrm{MC}(\Phi)$ for various classes Φ of first-order sentences. In this section we allow besides relation symbols also constants in the vocabulary of formulas in Φ. Abusing notation we continue to use notation like Σ_t and Π_t, now understanding that the formulas may contain constants.

Besides $|\varphi|$ it is technically convenient to also consider another size measure of φ, namely the number of subformulas $\|\varphi\|$ of φ counted with repetitions. In other words this is the number of nodes in the syntax tree of φ. Formally this is defined by recursion on syntax:

$$\begin{array}{rcll} \|\varphi\| & := & 1 & \text{for atomic } \varphi \\ \|(\varphi * \psi)\| & := & 1 + \|\varphi\| + \|\psi\| & \text{for } * \in \{\wedge, \vee\} \\ \|\neg\varphi\| & := & 1 + \|\varphi\| & \\ \|Qx\varphi\| & := & 1 + \|\varphi\| & \text{for } Q \in \{\forall, \exists\} \end{array}$$

Clearly we have $\|\varphi\| \leqslant |\varphi|$. Our model-checker will use as a "brute force" subroutine a straightforward model-checker which we describe first:

Lemma 6.1. *There is an algorithm deciding* $\mathrm{MC}(\mathrm{FO})$ *that on an instance* $(\varphi, \mathbf{A})$ *of* $\mathrm{MC}(\Sigma_1)$ *runs in space*

$$O\big(\|\varphi\| \cdot \log|A| + \log|\varphi| + \log|\mathbf{A}|\big).$$

Proof. We describe an algorithm $\mathbb{B}$ that decides the slightly more general (classical) problem

Instance:	a formula $\varphi = \varphi(\bar{x})$, a structure $\mathbf{A}$ and a tuple $\bar{a} \in A^{\lvert\bar{x}\rvert}$.
Problem:	$\mathbf{A} \models \varphi(\bar{a})$?

in space $O(\|\varphi\| \cdot \log|A| + \log n)$ on an instance $(\varphi, \mathbf{A}, \bar{a})$ of length n. The algorithm $\mathbb{B}$ implements a straightforward recursion on the logical syntax of φ.

If $\varphi(\bar{x}) = (\chi(\bar{x}) \wedge \psi(\bar{x}))$, it calls $\mathbb{B}(\chi(\bar{x}), \mathbf{A}, \bar{a})$. Upon returning from this call, $\mathbb{B}$ stores its answer, i.e. the bit b giving the truth value of $\mathbf{A} \models \chi(\bar{a})$. Then $\mathbb{B}$ calls $\mathbb{B}(\psi(\bar{x}), \mathbf{A}, \bar{a})$. Upon returning from this call with answer b', $\mathbb{B}$ answers the bit $b \cdot b'$.

If $\varphi(\bar{x}) = \exists x \chi(\bar{x}, x)$, then $\mathbb{B}$ loops over $b \in A$ and calls $\mathbb{B}(\psi(\bar{x}, x), \mathbf{A}, \bar{a}b)$; it answers with the maximal answer bit obtained.

The cases where φ is a negation, a disjunction or starts with $\forall$ are similarly explained.

If φ is atomic, it has the form $t_1{=}t_2$ or $R(t_1, \ldots, t_r)$ where R is an r-ary relation symbol and the t_i's are constants or variables. Assume the latter. Then $\mathbb{B}$ checks whether there exists $j \in [|\mathbf{A}|]$ such that for all $i \in [r]$ it holds that t_i^A equals the i-th component of the j-th tuple in $R^{\mathbf{A}}$. Here, t_i^A is $c^{\mathbf{A}}$ if t_i is a constant c; otherwise t_i is a free variable in φ and t_i^A is the corresponding component of $\bar{a}$.

To implement the recursion $\mathbb{B}$ stores a stack collecting the b's of the loops and the answer bits generated by the recursion as described. Scanning the whole stack allows to determine in space $O(\log n)$ the formula $\psi(\bar{x}, \bar{y})$ and tuple $\bar{b}$ such that thev corresponding recursive call is $\mathbb{B}(\psi(\bar{x}, \bar{y}), \mathbf{A}, \bar{a}\bar{b})$. The depth of the recursion is at most $\|\varphi\|$, so the stack can be stored in space $\|\varphi\| \cdot (\log|A| + 1)$. On an atomic formula $\mathbb{B}$ needs space $\log|\mathbf{A}| + \log|\varphi| \leqslant O(\log n)$ for the loops on j and i, and again $O(\log n)$ for the equality checks. Altogether we see that $\mathbb{B}$ can be implemented within the claimed space. □

The following proposition is the heart of the argument. The advantage with respect to the "brute force" algorithm from the previous proposition is that the factor $\|\varphi\|$ in the space bound is replaced by $\log\|\varphi\|$. But since other factors are worsened this algorithm is not in general more space efficient.

Proposition 6.2. *There is an algorithm deciding* $\mathrm{MC}(\Sigma_1)$ *that on an instance* $(\varphi, \mathbf{A})$ *of* $\mathrm{MC}(\Sigma_1)$ *runs in space*

$$O\big(\log\|\varphi\| \cdot w(\varphi) \cdot \log|A| + \log\|\varphi\| \cdot \log|\varphi| + \log|\mathbf{A}|\big).$$

Proof. We describe an algorithm $\mathbb{A}$ deciding the problem

Instance:	a Σ_1-formula φ, a natural number $w \geqslant w(\varphi)$, a structure $\mathbf{A}$ and $\bar{a} \in A^w$.
Problem:	$\mathbf{A} \models \varphi(\bar{a})$?

on an instance $(\varphi, w, \mathbf{A}, \bar{a})$ of length n in *allowed space*

$$O\big(\log\|\varphi\| \cdot w \cdot \log|A| + \log\|\varphi\| \cdot \log|\varphi| + \log n\big).$$

Notationally, $\mathbf{A} \models \varphi(\bar{a})$ means that the assignment that maps the i-th largest free variable in φ to the i-th component of $\bar{a}$ satisfies φ in $\mathbf{A}$. Here we suppose an order of the variables in φ, say

according to appearance in φ. All we need is that we can determine in space $O(\log n)$ the value assigned to a given variable in a given a subformula of φ according to a given tuple $\bar{a} \in A^w$.

For a sufficiently large constant $c \in \mathbb{N}$ to be determined in the course of proof, $\mathbb{A}$ checks that

$$\|\varphi\| \geqslant c \cdot w + c. \tag{16}$$

If this is not the case, then $\mathbb{A}$ uses "brute force", that is, runs the algorithm from Lemma 6.1.

Now suppose (16) holds. Choosing $c \geqslant 3$ this implies that the syntax tree of φ has at least 3 nodes. Then $\mathbb{A}$ computes in space $O(\log |\varphi|)$ a subformula φ_0 of φ such that

$$\|\varphi\|/3 \leqslant \|\varphi_0\| \leqslant 2\|\varphi\|/3. \tag{17}$$

Observe that φ_0 is a Σ_1-formula. Let $\bar{y} = y_1 \cdots y_{|\bar{y}|}$ list the free variables of φ_0 and note

$$|\bar{y}| \leqslant w(\varphi_0) \leqslant w(\varphi) \leqslant w. \tag{18}$$

Let $c_1, \ldots, c_{|\bar{y}|}$ be new constant symbols. For every free variable y_i of φ_0 check whether it has an occurrence in φ_0 which is not a free occurrence in φ. If such an occurrence exists, all the free occurrences of y_i in φ_0 appear within a uniquely determined subformula $\exists y_i \chi$ of φ containing φ_0 as a subformula, where $\exists y_i$ binds these occurrences, that is, the free occurrences of y_i in φ_0 are also free in χ. Replace $\exists y_i \chi$ by $\exists y_i (y_i{=}c_i \wedge \chi)$. Let φ_1 denote the resulting formula. Note that φ_1 does not depend on the order of how these replacements are performed.

Moreover, using (18),

$$\big\|\varphi_1\big\| \leqslant \|\varphi\| + 2|\bar{y}| \leqslant \|\varphi\| + 2w. \tag{19}$$

The algorithm $\mathbb{A}$ then loops through $\bar{b} = (b_1, \ldots, b_{|\bar{y}|}) \in A^{|\bar{y}|}$ and does two recursive calls:

(R0) Recursively call $\mathbb{A}(\varphi_0, w, \mathbf{A}, \bar{b})$ to check whether $\mathbf{A} \models \varphi_0(\bar{b})$. If $\mathbf{A} \models \varphi_0(\bar{b})$, then replace φ_0 in φ_1 by $c_1{=}c_1$; otherwise by $\neg c_1{=}c_1$.

Let $\varphi_1^{\bar{b}}$ be the resulting formula. Further, let $\mathbf{A}^{\bar{b}}$ be the expansion of $\mathbf{A}$ that interprets the constants $c_1, \ldots, c_{|\bar{y}|}$ by $b_1, \ldots, b_{|\bar{y}|}$ respectively.

(R1) Recursively call $\mathbb{A}\big(\varphi_1^{\bar{b}}, w, \mathbf{A}^{\bar{b}}, \bar{a}\big)$ and output its answer.

Note that in (R1) we have $w(\varphi_1^{\bar{b}}) \leqslant w(\varphi) \leqslant w$, so the algorithm recurses to an instance of our problem. It is routine to verify that

$$\mathbf{A} \models \varphi(\bar{a}) \quad \Longleftrightarrow \quad \text{there exists } \bar{b} \in A^s \text{ such that } \mathbf{A}^{\bar{b}} \models \varphi^{\bar{b}}(\bar{a}).$$

Thus, $\mathbb{A}$ correctly decides whether $\mathbf{A} \models \varphi(\bar{a})$. It remains to show that $\mathbb{A}$ can be implemented in the allowed space.

We first estimate the depth of the recursion. In (R0) the algorithm recurses to formula φ_0 and $\|\varphi_0\| \leqslant 2\|\varphi\|/3$ by (17). In (R1) the algorithm recurses on $\varphi_1^{\bar{b}}$, a formula obtained from φ_1 by replacing the subformula φ_0 by an atomic formula or the negation of an atomic formula. Then

$$\|\varphi_1^{\bar{b}}\| \leqslant \|\varphi\| + 2w - \|\varphi_0\| + 1 \leqslant 3\|\varphi\|/4 + (2w + 1 - \|\varphi\|/12)$$

by (19) and (17). Provided c in (16) is large enough, this implies $\|\varphi_1^{\bar{b}}\| \leqslant 3\|\varphi\|/4$. It follows that the recursion depth is $O(\log \|\varphi\|)$.

In each recursive call the structure the algorithm recurses to is determined by a tuple $\bar{b}$ from the loop. This tuple has length $\leqslant w$ (cf. (18)). The formula $\varphi_1^{\bar{b}}$ in (R1) is determined by the truth value of $\mathbf{A} \models \varphi_0(\bar{b})$. The algorithm recurses either

(P0) to the formula φ_0, or

(P1) to the formula $\varphi_1^{\bar{b}}$ as defined if $\mathbf{A} \models \varphi_0(\bar{b})$, or

(P2) to the formula $\varphi_1^{\bar{b}}$ as defined if $\mathbf{A} \not\models \varphi_0(\bar{b})$.

To implement the recursion, $\mathbb{A}$ maintains a sequence of tuples $\bar{b}_1, \dots, \bar{b}_d$ and a sequence of "possibilities" $(p_1, \dots, p_d) \in \{0,1,2\}^d$. The length d is bounded by $O(\log \|\varphi\|)$, the depth of the recursion. In a recursive call as described above the sequence of tuples is expanded by the tuple $\bar{b}$ from the loop, and the possibility sequence by $0, 1, 2$ depending on which recursion from (P0), (P1), (P2) is taken. Both sequences can be stored in allowed space, namely

$$O(d \cdot w \cdot \log |A|) \leqslant O(\log \|\varphi\| \cdot w \cdot \log |A|).$$

We still have to explain how, given $(\bar{b}_1, \dots, \bar{b}_d)$ and $(p_1, \dots, p_d)$, the algorithm determines the corresponding formula ψ and structure $\mathbf{B}$ it has to recurse to. The structure $\mathbf{B}$ is an expansion of the input structure $\mathbf{A}$ by a sequence of constants interpreted by the sequence $\bar{b}_{i_1} \cdots \bar{b}_{i_\ell}$ where $i_1 < \cdots < i_\ell$ list those indices $j \in [d]$ such that $p_j \neq 0$. This structure can be computed from the two sequences and the input in logarithmic space. Note that $\mathbf{B}$ is an expansion of $\mathbf{A}$ by interpreting at most $w \cdot d$ many constants, so

$$\log |\mathbf{B}| \leqslant \log(|\mathbf{A}| + 2 \cdot w \cdot d) \leqslant \log |\mathbf{A}| + \log w + \log\log \|\varphi\| + O(1). \tag{20}$$

To determine the formula ψ, consider the function R that maps a formula χ and a number $i \leqslant 2$ to the formula according (Pi). Since we defined the recursion possibilities (Pi) only on formulas satisfying (16), let us agree that $R(\chi, i) := \chi$ on formulas violating (16). The desired formula ψ is obtained by iterating this function along $(p_1, \dots, p_d)$: compute $\psi_1 := R(\varphi, p_1)$, $\psi_2 := R(\psi_1, p_2)$, …and output $\psi := \psi_d$. Note that all these formulas have length $O(|\varphi|)$ and each iteration step is computable in space $O(\log |\varphi|)$. Hence, the whole iteration can be implemented in space

$$O(d \cdot \log |\varphi|) \leqslant O(\log \|\varphi\| \cdot \log |\varphi|).$$

If $(\psi, w, \mathbf{B}, \bar{b})$ is such that $\|\psi\|$ violates (16), i.e. $\|\psi\| < c \cdot w + c$, then $\mathbb{A}$ invokes the "brute force" algorithm from Lemma 6.1. This requires space

$$O\big(w \cdot \log |A| + \log |\psi| + \log |\mathbf{B}|\big).$$

By $|\psi| \leqslant O(|\varphi|)$ and (20), this is allowed space. □

We now extend the space bound from the previous proposition to $\mathrm{MC}(\Sigma_t)$ for each $t \geqslant 1$. Recalling $\|\varphi\| \leqslant |\varphi|$, the result implies Theorem 1.3.

Theorem 6.3. *For all $t \geqslant 1$ there is an algorithm deciding $\mathrm{MC}(\Sigma_t)$ that on an instance $(\varphi, \mathbf{A})$ of $\mathrm{MC}(\Sigma_t)$ runs in space*

$$O\big(\log\|\varphi\| \cdot w(\varphi) \cdot \log|A| + \log\|\varphi\| \cdot \log|\varphi| + \log|\mathbf{A}|\big).$$

Proof. Similarly as in the previous proposition, we give an algorithm $\mathbb{A}$ deciding

Instance:	a Σ_t-formula φ, a natural number $w \geqslant w(\varphi)$, a structure $\mathbf{A}$ and $\bar{a} \in A^w$.
Problem:	$\mathbf{A} \models \varphi(\bar{a})$?

on an instance $(\varphi, w, \mathbf{A}, \bar{a})$ of length n in *allowed space*

$$O\big(\log\|\varphi\| \cdot w \cdot \log|A| + \log\|\varphi\| \cdot \log|\varphi| + \log n\big).$$

For $t = 1$ this is what has been shown in the proof of Proposition 6.2. So we assume $t \geqslant 2$ and proceed inductively.

Let $\psi_1, \dots, \psi_r \in \Pi_{t-1}$ be such that φ results from these formulas by existential quantification and positive Boolean combinations. Of course, such formulas ψ_j are computable in logarithmic space from φ. For $j \in [r]$ let $s_j \leqslant w(\varphi) \leqslant w$ be the number of variables occurring freely in ψ_j and let $\bar{x}_j$ be a length s_j tuple listing these variables.

Let $\mathbf{A}^*$ be the structure with universe A that interprets for each $j \in [r]$ an s_j-ary relation symbol R_j by

$$R_j^{\mathbf{A}^*} := \{\bar{b} \in A^{s_j} \mid \mathbf{A} \models \psi_j(\bar{b})\}.$$

By the induction hypothesis we can compute $\mathbf{A}^*$ in allowed space. Moreover

$$|\mathbf{A}^*| \leqslant \|\varphi\| + |A| + \|\varphi\| \cdot |A|^w. \tag{21}$$

Define the formula φ^* by replacing for every $j \in [r]$ the formula $\psi_j(\bar{x})$ in φ by $R_j(\bar{x}_j)$. Clearly, we have $\varphi^* \in \Sigma_1^s$ and

$$\|\varphi^*\| \leqslant \|\varphi\| \quad \text{and} \quad |\varphi^*| \leqslant O(|\varphi|). \tag{22}$$

More importantly,

$$\mathbf{A} \models \varphi(\bar{a}) \iff \mathbf{A}^* \models \varphi^*(\bar{a}).$$

Since $w(\varphi^*) \leqslant w(\varphi) \leqslant w$ we have with $(\varphi^*, w, \mathbf{A}^*, \bar{a})$ an instance of $\mathrm{MC}(\Sigma_1)$. By Proposition 6.2, we can decide whether $\mathbf{A}^* \models \varphi^*(\bar{a})$ in space

$$O\big(\log\|\varphi^*\| \cdot w \cdot \log|A| + \log\|\varphi^*\| \cdot \log|\varphi^*| + \log|\mathbf{A}^*|\big)$$

By (22) and (21), this is allowed space. □

7 Future directions

While this paper focusses on the model-checking problem for bounded variable first-order logic, it points to some more general problems and directions of independent interest.

As mentioned in the introduction, we view the results about the tree hierarchy as a contribution to a fine-structure theory of FPT. Recent works [14, 1, 9] on parameterized circuit complexity also go in this direction. Yet missing is a theory of FPT-completeness with respect to pl-reductions. Corollary 4.5 indicates that such a theory is not entirely analogous to the classical theory of P-completeness.

We iterate the question whether PATH or TREE are closed under complementation. The question begs for a better understanding of inductive counting [2] and its limits. We noted that a positive answer for TREE would imply a collapse of the tree hierarchy (Corollary 4.2). We do not know whether this also follows from PATH being closed under complementation.

As a further structural question we do not know how para-NL relates to the tree hierarchy. Proposition 3.4 (1) gives only a partial answer.

We showed that the straightforward model-checking algorithm is space-optimal under the hypothesis that Savitch's theorem is optimal (Theorem 1.2). We conjecture that similar optimality results can be derived for other natural algorithms as well. Future work will show to what extent the hypothesis that Savitch's theorem is optimal can play a role in space complexity similar to the one played by the ETH in time complexity.

References

[1] M. Bannach, C. Stockhusen, and T. Tantau. Fast parallel fixed-parameter algorithms via color coding. *Proceedings of IPEC 2015*, LIPIcs 43, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 224?235, 2015.

[2] A. Borodin, S. A. Cook, P. W. Dymond, W. L. Ruzzo, and M. Tompa. Two applications of inductive counting for complementation problems. *SIAM Journal on Computing* 18(3): 559–578, 1989.

[3] L. Cai, J. Chen, R. G. Downey, and M. R. Fellows. On the structure of parameterized problems in NP. *Information and Computation* 123: 38–49, 1995.

[4] L. Cai, J. Chen, R. G. Downey, and M. R. Fellows. Advice classes of parameterized tractability. *Annals of Pure and Applied Logic* 84(1): 119–138, 1997.

[5] H. Chen. On the complexity of existential positive queries. *ACM Transactions on Computational Logic* 15 (1), 2014.

[6] H. Chen and M. Müller. One hierarchy spawns another: graph deconstructions and the complexity classification of conjunctive queries. *Proceedings of CSL-LICS'14*, pp. 32:1-32:10, 2014. Full version available as arXiv:1307.1353 [cs.CC].

[7] H. Chen and M. Müller. The fine classification of conjunctive queries and parameterized logarithmic space complexity. *ACM Transactions on Computation Theory* 7 (2): Article No. 7, 2015.

[8] Y. Chen and J. Flum. On optimal inverters. *The Bulletin of Symbolic Logic* 20 (01), 1-23, 2014.

[9] Y. Chen and J. Flum. Some lower bounds in parameterized AC0. *Proceedings of MFCS'16*, LIPIcs 58, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 27:1-27:14, 2016.

[10] Y. Chen, J. Flum and M. Grohe. Machine-based methods in parameterized complexity. *Theoretical Computer Science* 339 (2-3): 167-199, 2005.

[11] Y. Chen and M. Müller. Bounded variable logic, parameterized logarithmic space and Savitch's theorem. *Proceedings of MFCS'14*, LNCS 8634, pp. 183-195, 2014.

[12] V. Dalmau, P. G. Kolaitis and M. Y. Vardi. Constraint Satisfaction, Bounded Treewidth, and Finite-Variable Logics. *Proceedings of CP'02*, LNCS 2470, pp. 310–326, 2002.

[13] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. 2nd ed., Springer Monographs in Mathematics 2006, 1999.

[14] M. Elberfeld, C. Stockhusen, and T. Tantau. On the space and circuit complexity of parameterized problems: classes and completeness. *Algorithmica* 71(3): 661–701, 2015.

[15] J. Flum and M. Grohe. Describing parameterized complexity classes. *Information and Computation*, 187(2):291–319, 2003.

[16] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.

[17] J. Flum and M. Grohe. Fixed-parameter tractability, definability, and model checking. *SIAM Journal on Computing* 31:113–145, 2001.

[18] J. Flum, M. Grohe and M. Weyer. Bounded fixed-parameter tractability and $\log^2 n$ nondeterministic bits , *Journal of Computer and System Sciences* 72: 34–71, 2006.

[19] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48(6):1184–1206, 2001.

[20] J. Gao and R. Impagliazzo. Orthogonal Vectors is hard for first-order properties on sparse graphs. *ECCC*, TR16-053, Revision #1, 2016.

[21] M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1):1:1–1:24, 2007.

[22] M. Grohe. Logic, graphs, and algorithms. In J. Flum, E. Grädel, T. Wilke (Eds), *Logic and Automata – History and Perspectives*, Amsterdam University Press, 2007.

[23] M. Grohe. Algorithmic meta theorems for sparse graph classes. *Proceedings of CSR'14*, LNCS 8476, pp.16-22, 2014.

[24] M. Grohe and S. Kreutzer. Methods for algorithmic meta theorems. In *Model Theoretic Methods in Finite Combinatorics*, Contemporary Mathematics 558, AMS, 2011.

[25] M. Grohe, T. Schwentick, and L. Segoufin. When is the evaluation of conjunctive queries tractable? *Proceedings of STOC'01*, pp. 657–666, 2001.

[26] L. A. Hemaspaandra, M. Ogihara, S. Toda. Space-efficient recognition of sparse self-reducible languages. *Computational Complexity* 4: 262–296, 1994.

[27] N. Immerman. Nondeterministic space is closed under complementation. *SIAM Journal on Computing*, 17(5): 935–938, 1988.

[28] P. G. Kolaitis and M. Y. Vardi. Conjunctive-query containment and constraint satisfaction. *Journal of Computer and System Sciences*, pp. 302–332, 2000.

[29] L. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3): 265–266, 1973.

[30] J. R. Lipton. *The P = NP Question and Gödel's Lost Letter*. Springer, 2010.

[31] J.-A. Montoya and M. Müller. Parameterized random complexity. *Theory of Computing Systems* 52(2): 221–270, 2013.

[32] M. Müller. Parameterized logarithmic space and fine structure of FPT. *FPT News: The Parameterized Complexity Newsletter Vol. 10 No. 2*, September 2014.

[33] A. Potechin. Bounds on monotone switching networks for directed connectivity. In *Proceedings of FOCS'10*, pp. 553–562, 2010.

[34] W. L. Ruzzo. Tree-size bounded alternation. *Journal of Computer and System Sciences* 21(2): 218–235, 1980.

[35] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2): 177–192, 1970.

[36] N. Schweikardt, T. Schwentick, and L. Segoufin. Database theory: Query languages. *Algorithms and theory of computation*, 19.1–19.34. Chapman & Hall/CRC, 2010.

[37] R. Szelepcsényi. The method of forcing for nondeterministic automata. *Bulletin of the EATCS* 33:96–100, 1987.

[38] M. Y. Vardi. On the complexity of bounded-variable queries. *Proceedings of PODS'95*, pp. 266–276. ACM Press, 1995.

[39] H. Venkateswaran. Properties that characterize LOGCFL. *Journal of Computer and System Sciences* 43(2): 380–404, 1991.

[40] A. Wigderson. The complexity of graph connectivity. *Proceedings of MFCS'92*, LNCS 629, pp. 112–132, 1992.

[41] R. Williams. Faster decision of first-order graph properties. *Proceedings of CSL-LICS'14*, Article No. 80, 2014.